




# Kubernetes basics





# Kubernetes basics



1. Cluster
2. Namespaces
3. Working with a cluster
4. Everything runs on Nodes
5. Pod
6. Service / Deployment
7. Horizontal Pod Autoscaler
8. Deployment strategy
9. Ingress
10. Persistent Volume Claim
11. Configmap
12. Statefullset
13. Job
14. Cron
15. Loadtesting
16. Monitoring
17. Logs

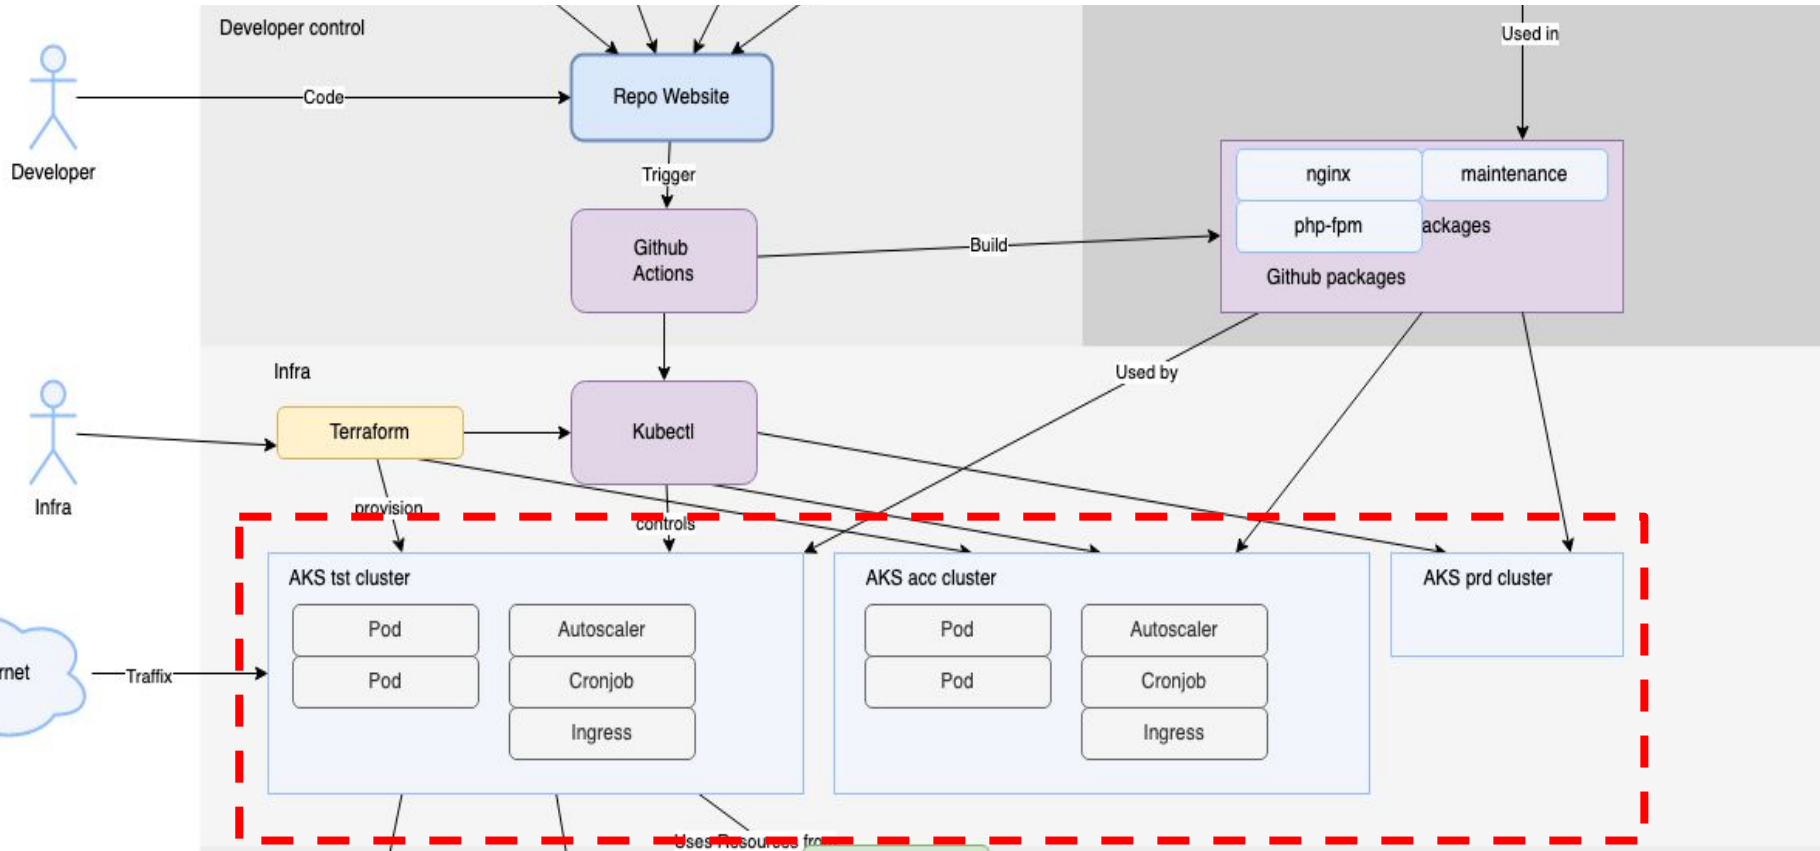


Blogpost  
with some c  
this conten





# Clusters

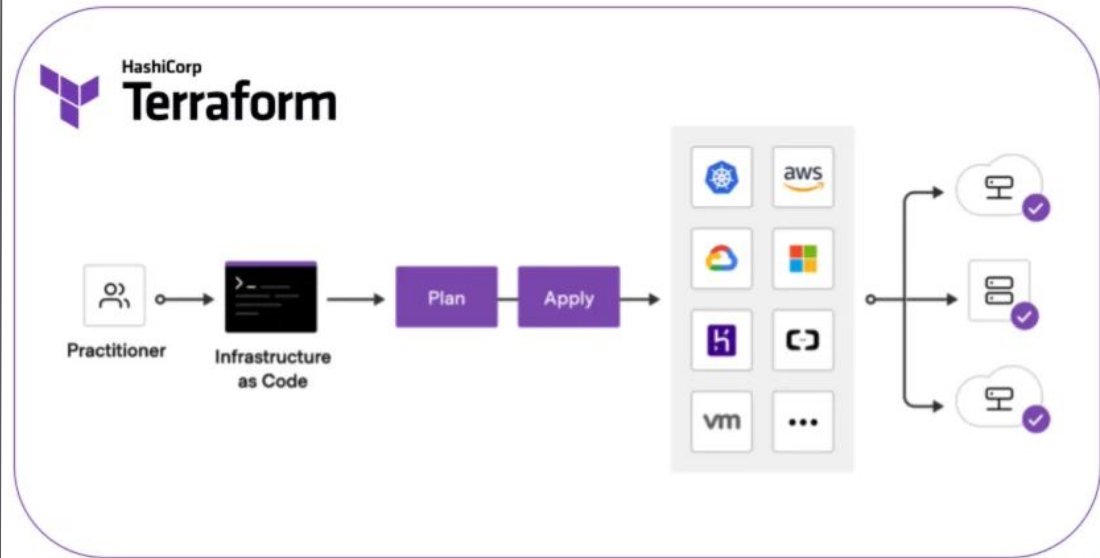




# Cluster creation

Manual


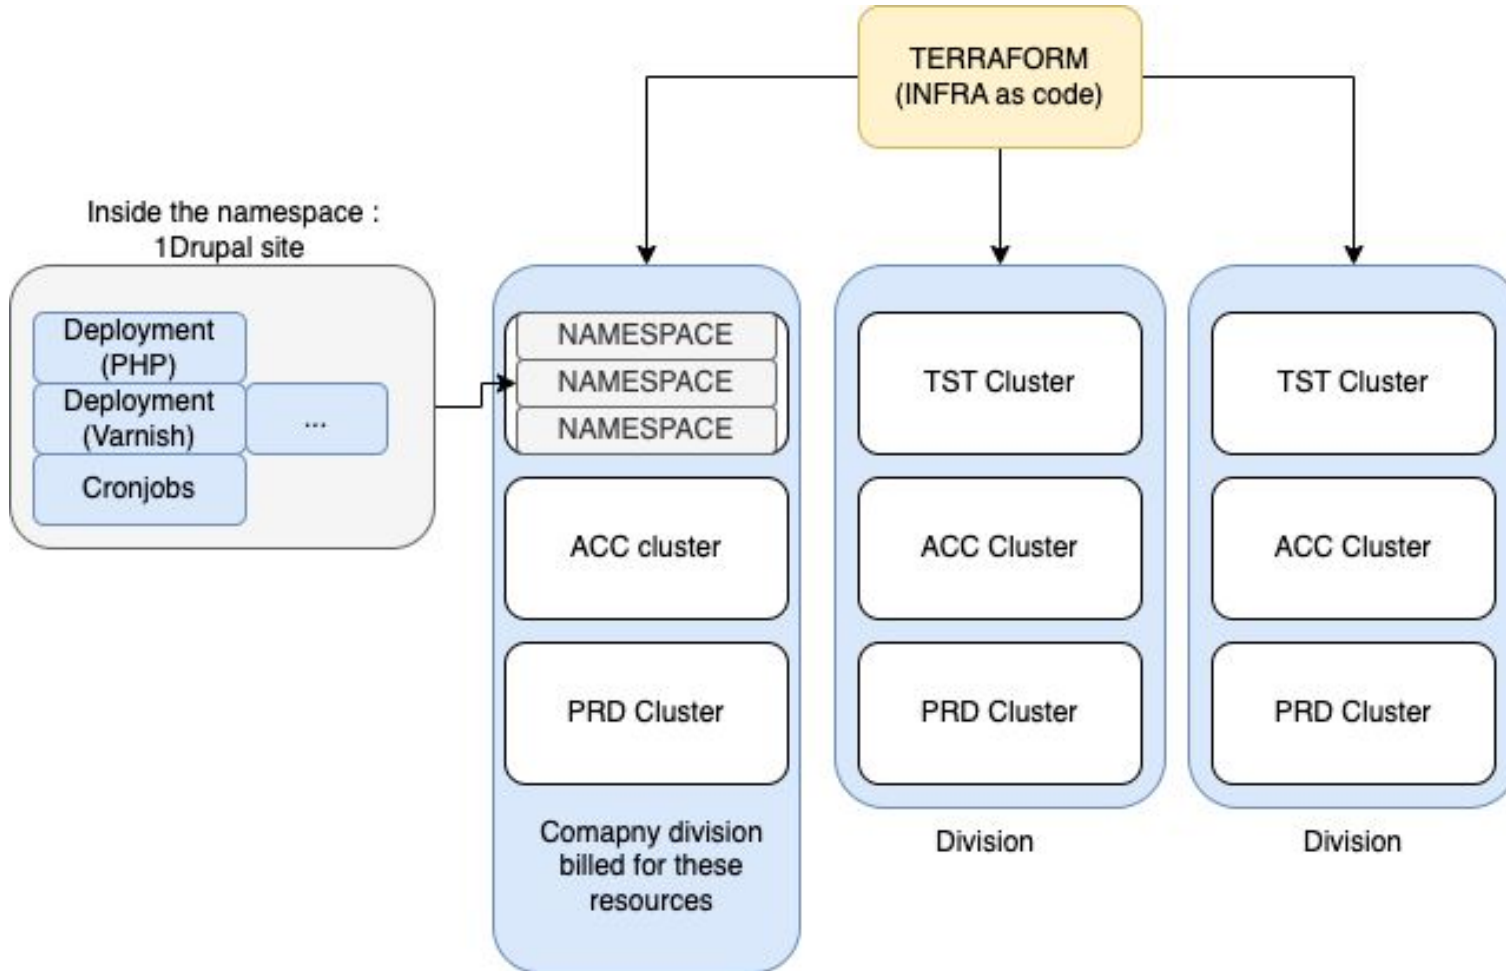
in Code






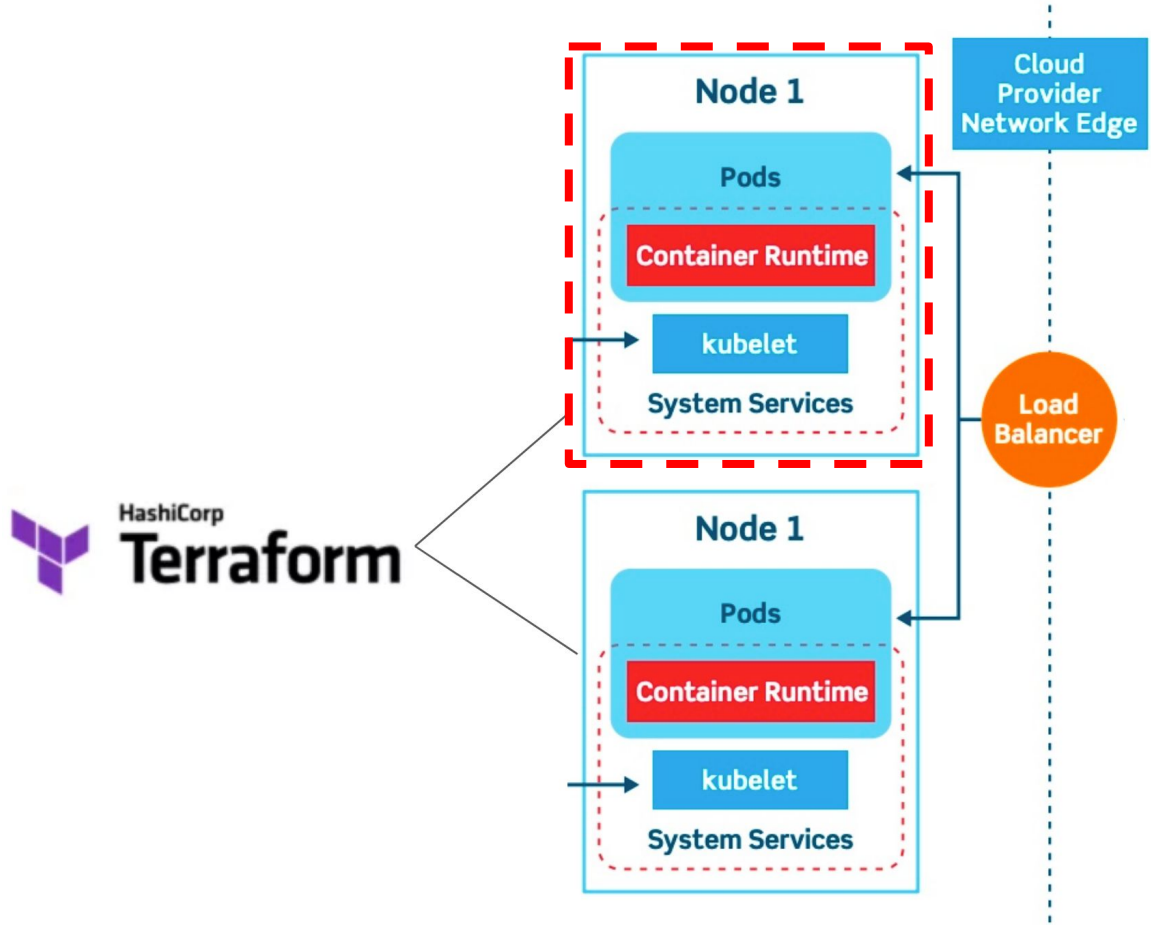


# Namespaces





# Everything runs on Nodes

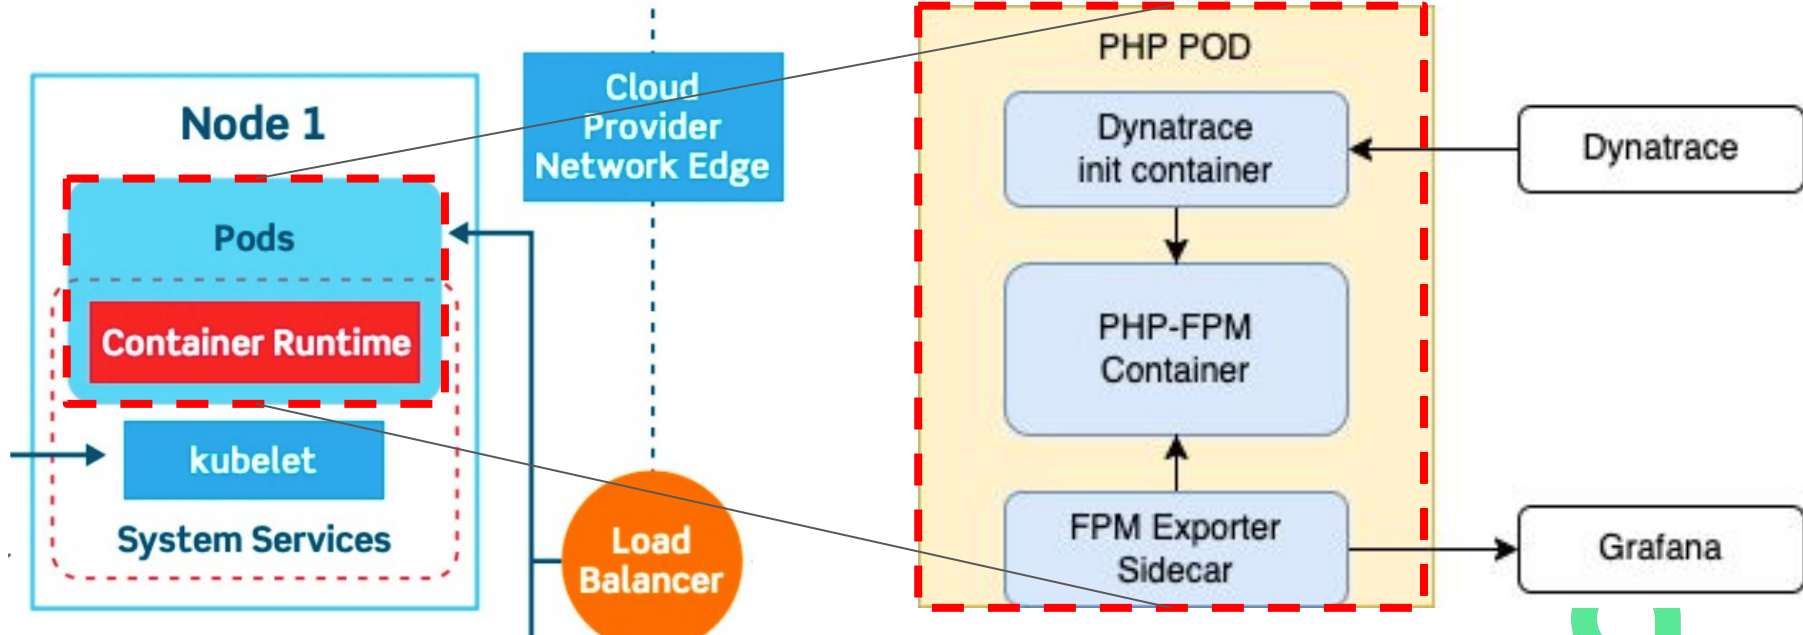




# Pod

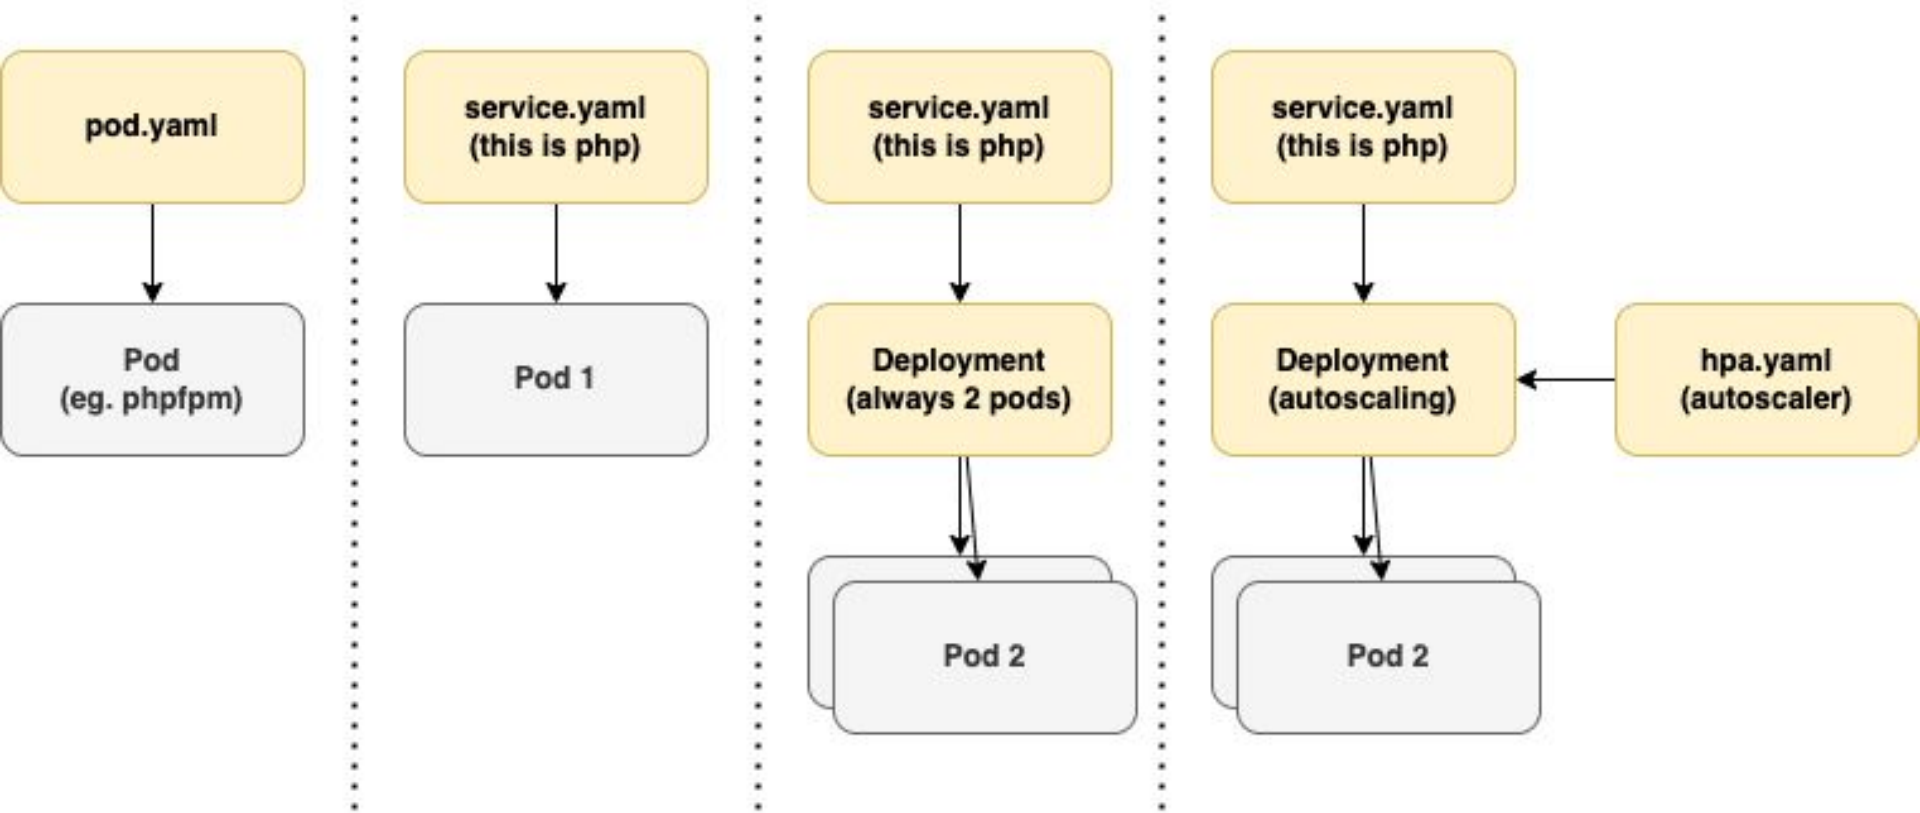


## Real life example:



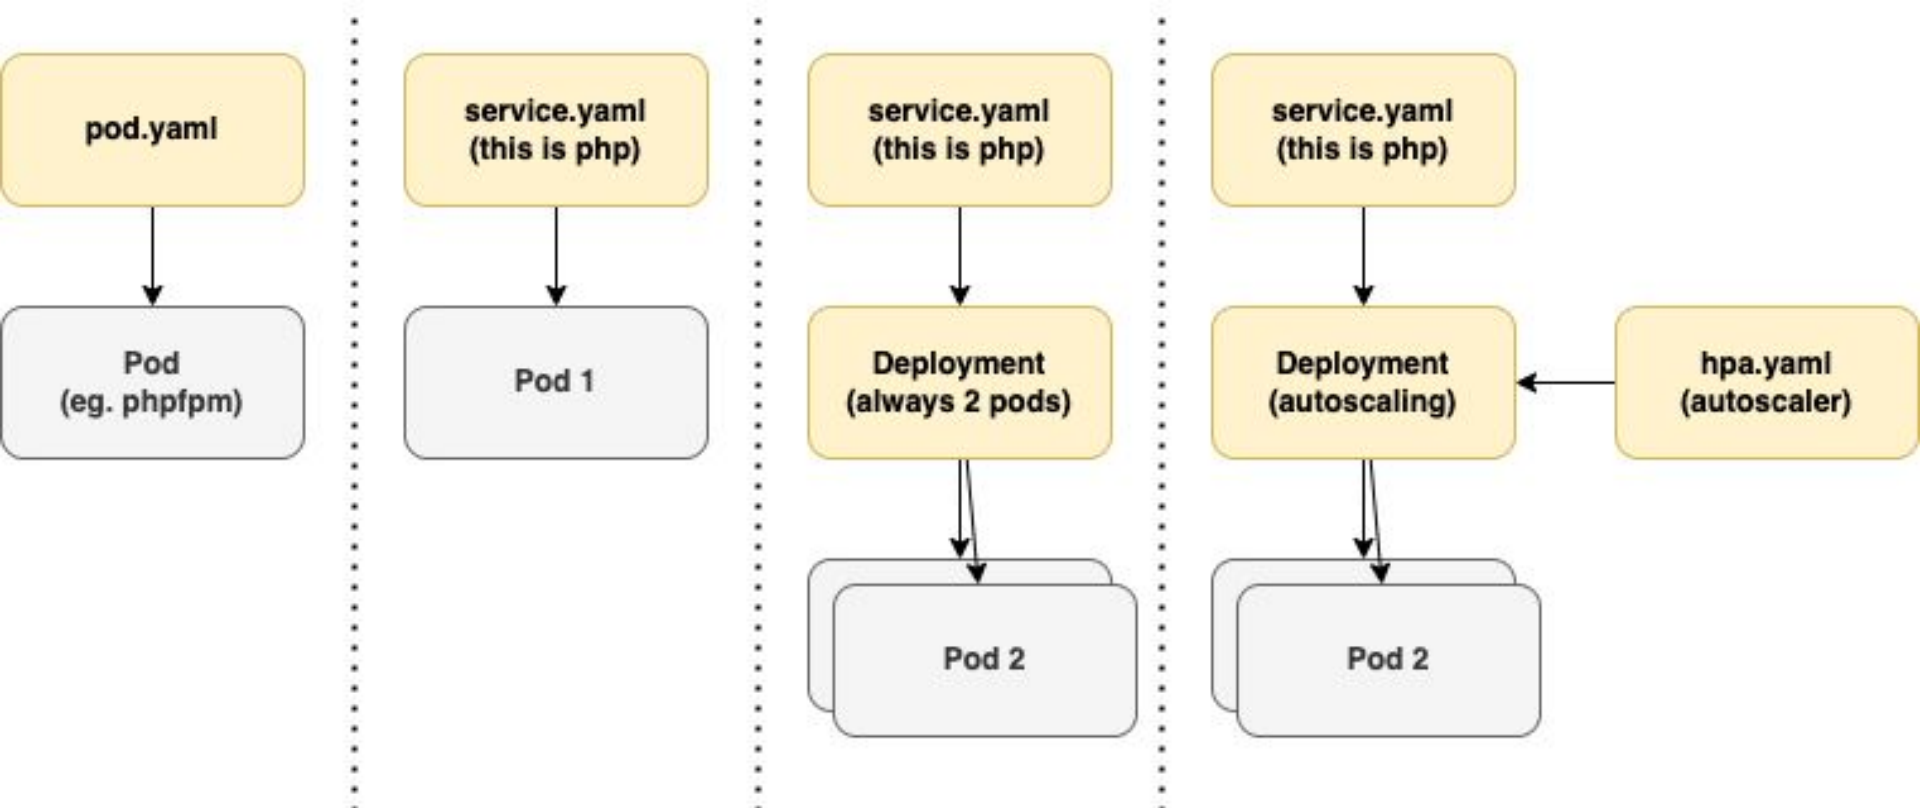


# Pod - Service - Deployment - Autoscaling





# Pod - Service - Deployment - Autoscaling



**NEED TO KNOW:** HPA needs resource limitations set on pods & CPU < 1 will throttle your system.





# Kubernetes - autoscaling

Horizontal Pod Autoscalers 1 item

Namespace: orderentryaccountants-t

Search Horizontal Pod Autoscalers

<input type="checkbox"/>	Name	Namesp...	Metrics	Min Pods	Max Pods	Replicas	A...	Status	
<input type="checkbox"/>		orderentrya...	1% / 70%	1	10	1	85d	AbleToScale	Scaling

Defined the pods as a deployment  
With resources

```
resources:  
  requests:  
    cpu: 100m  
    memory: "256Mi"  
  limits:  
    cpu: 1  
    memory: "4112Mi"
```

Describe the  
HPA




```
apiVersion: autoscaling/v2beta2  
kind: HorizontalPodAutoscaler  
metadata:  
  name: {{.Release.Name}}-php-fpm-hpa  
  namespace: {{.Release.Namespace}}  
spec:  
  scaleTargetRef:  
    apiVersion: apps/v1  
    kind: Deployment  
    name: {{.Release.Name}}-php-fpm  
  minReplicas: 1  
  maxReplicas: 10  
  metrics:  
  - type: Resource  
    resource:  
      name: cpu  
      target:  
        type: Utilization  
        averageUtilization: 70
```



# Kubernetes - Deployment strategy


✨ more cool stuff ✨

- **recreate**: terminate the old version and release the new one
- **ramped**: release a new version on a rolling update fashion, one after the other
- **blue/green**: release a new version alongside the old version then switch traffic
- **canary**: release a new version to a subset of users, then proceed to a full rollout
- **a/b testing**: release a new version to a subset of users in a precise way (HTTP headers, cookie, weight, etc.). This doesn't come out of the box with Kubernetes, it imply extra work to setup a smarter loadbalancing system (Istio, Linkerd, Traeffik, custom nginx/haproxy, etc).
- **shadow**: release a new version alongside the old version. Incoming traffic is mirrored to the new version and doesn't impact the response.







# Drupal & Zero downtime deploys



No Schema updates



Schema updates





# Drupal Zero downtime deploys

D03 Deploy to PRD  
deploy-prd.yml

Filter workflow runs

33 workflow runs

Event - Status - Branch - Actor -

This workflow has a workflow\_dispatch event trigger. Run workflow -

**D03 Deploy to PRD**  
D03 Deploy to PRD #33: Manually run by lammensj-ds

**D03 Deploy to PRD**  
D03 Deploy to PRD #32: Manually run by lammensj-ds

**D03 Deploy to PRD**  
D03 Deploy to PRD #31: Manually run by lammensj-ds

**D03 Deploy to PRD**  
D03 Deploy to PRD #30: Manually run by lammensj-ds

**D03 Deploy to PRD**  
D03 Deploy to PRD #29: Manually run by lammensj-ds

2 weeks ago ...  
10m 13s

2 weeks ago ...  
12m 53s

Use workflow from

Branch: master -

Git tag to deploy (eg 0.0.0) +

Deployment from +

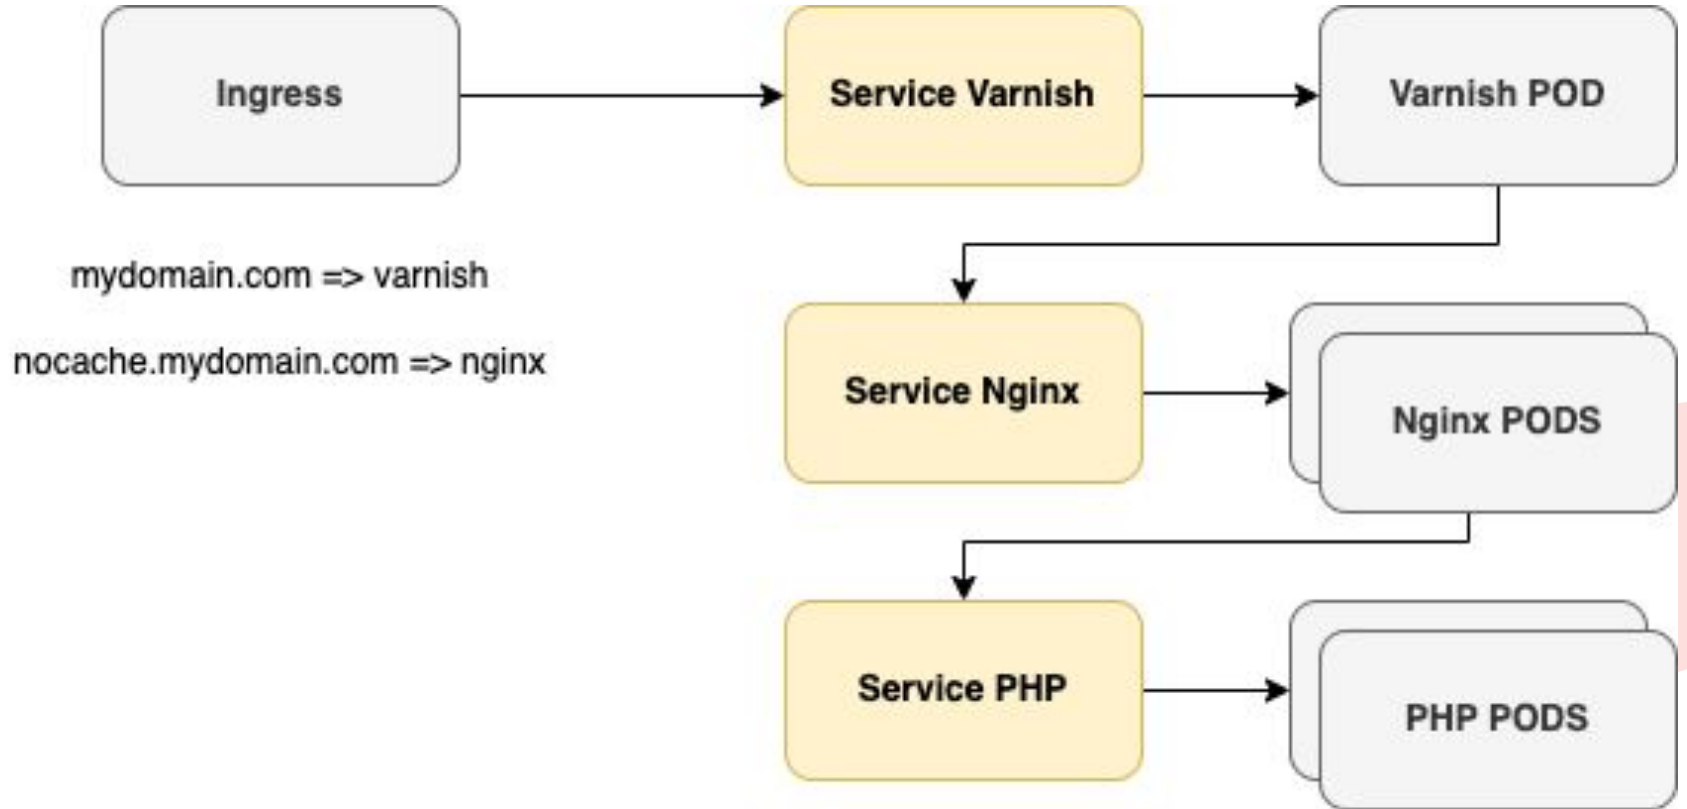
rolling-deploy

✓ schema-deploy

Run workflow

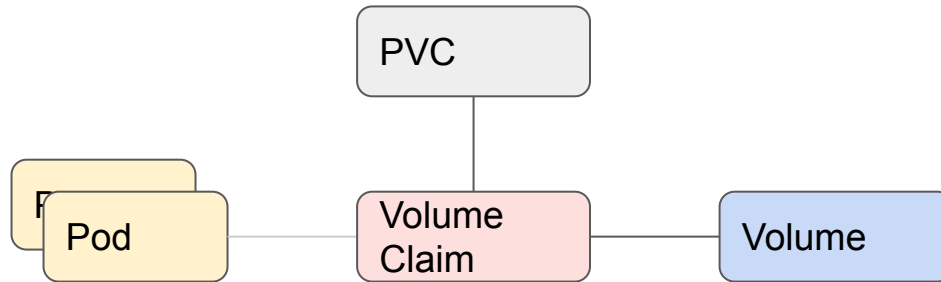


# Ingress





# Persistent Volume Claim



eg:  
/sites/default/files  
/tmp  
/private




# Configmap

## MAIN USE

Especially practical if you want to use an

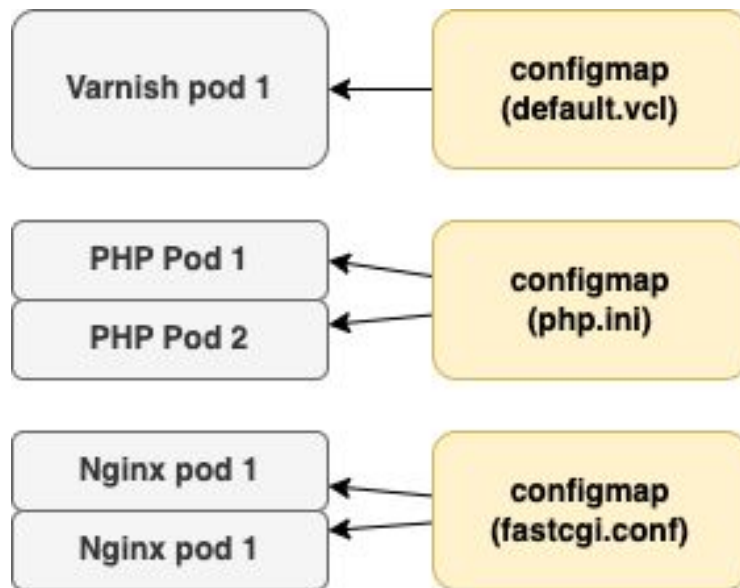
**external docker container** image (eg memcache or varnish)

with your specific config (without creating your own images)





# Configmap



## PRACTICAL FOR DEBUGGING

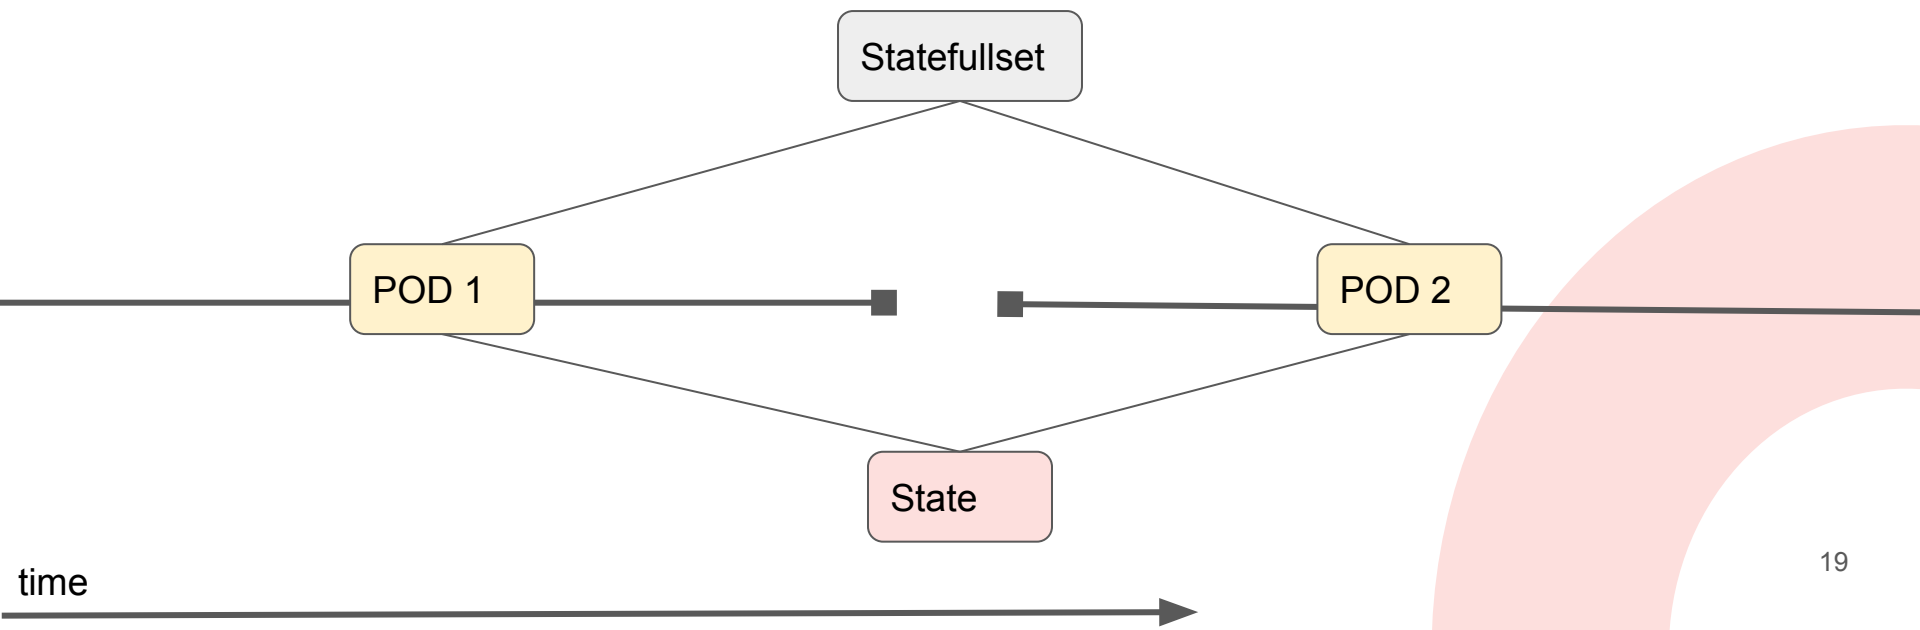
Also super practical to add php.ini or nginx or apache config to your custom images.

This adds the possibility to fiddle on an (obviously not production) environment with the config without a full build + deployment (which is slow).

Change configmap + delete pod (triggers recreate with new config)



# StatefullSet





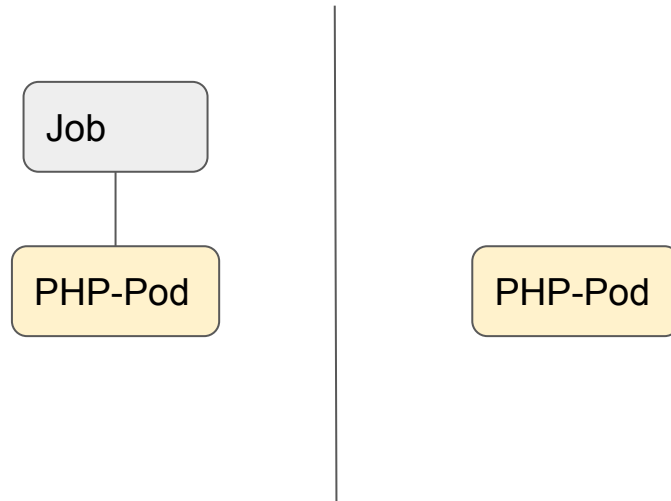
# StatefulSet

eg: memcache or redis

```
---
apiVersion: apps/v1
kind: StatefulSet
metadata:
  name: {{.Release.Name}}-memcached
  namespace: {{.Release.Namespace}}
spec:
  selector:
    matchLabels:
      app: {{.Release.Name}}-memcached
  serviceName: {{.Release.Name}}-memcached
  replicas: 1
  template:
    metadata:
      labels:
        app: {{.Release.Name}}-memcached
    spec:
      containers:
        - name: {{.Release.Name}}-memcached
          image: memcached
          args: ["-m 64 -p 11211 -p 0.0.0.0"]
          ports:
            - containerPort: [REDACTED]
              protocol: TCP
          readinessProbe:
            tcpSocket:
              port: [REDACTED]
            initialDelaySeconds: 5
            periodSeconds: 10
```

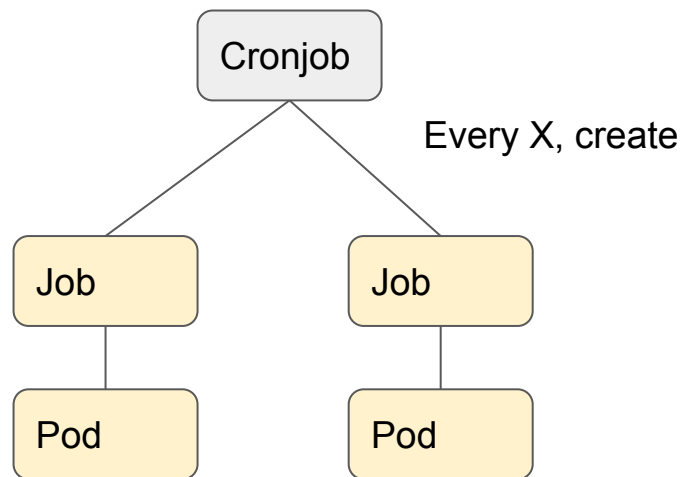


# Job





# Cron





# Cron

```
---
apiVersion: batch/v1
kind: CronJob
metadata:
  name: "{{.Release.Name}}-cron"
  namespace: "{{.Release.Namespace}}"
spec:
  suspend: true
  schedule: "*/5 * * * *"
  concurrencyPolicy: Forbid
  jobTemplate:
    spec:
      template:
        spec:
```

```
spec:
  template:
    spec:
      {{- with .Values.imagePullSecrets }}
      imagePullSecrets:
        {{- toYaml . | nindent 12 }}
      {{- end }}
      volumes:
      - name: files-vol
        persistentVolumeClaim:
          claimName: "{{.Release.Name}}"
      - name: private-files-vol
        persistentVolumeClaim:
          claimName: "{{.Release.Name}}-private"
      - name: tmp-files-vol
        persistentVolumeClaim:
          claimName: "{{.Release.Name}}-tmp"
      containers:
      - name: drupal-cron
        image: "{{.Values.php_image_name}}:{{.Values.image.version}}"
        imagePullPolicy: IfNotPresent
        volumeMounts:
        - mountPath: /var/www/drupal/docroot/sites/default/files
```

# Loadtesting (job)

```
apiVersion: batch/v1
kind: CronJob
metadata:
  name: "{{.Release.Name}}-loadtest"
  namespace: "{{.Release.Namespace}}"
spec:
  suspend: true
  schedule: "1 1 * * 1"
  concurrencyPolicy: Forbid
  jobTemplate:
    spec:
      template:
        spec:
          volumes:
            - name: k6-loadtest
              configMap:
                name: k6-loadtest
                items:
                  - key: test.js
                    path: test.js
          containers:
            - name: "{{.Release.Name}}-k6"
              image: loadimpact/k6:latest
              imagePullPolicy: Always
              env:
                - name: HOSTNAME
                  value: "{{.Values.ingress_hostname }}"
              command: ["k6"]
              args: ["run", "/test.js"]
```



```
└─ etc
  └─ drupal
  └─ drush
  └─ git
  └─ k6
    └─ JS test.js
```



# Monitoring

<input type="checkbox"/>	...		0	ReplicaSet	aks-nptstiws:	BestEffort	34d	Running
<input type="checkbox"/>	...		0	StatefulSet	aks-nptstiws:	BestEffort	34d	Running
<input type="checkbox"/>	n...		0	ReplicaSet	aks-nptstiws:	BestEffort	18h	Running
<input type="checkbox"/>	p...		0	ReplicaSet	aks-nptstiws:	Burstable	18h	Running
<input type="checkbox"/>	v...		0	ReplicaSet	aks-nptstiws:	BestEffort	27d	Running



⚠ ⚠ ⚠ Sidecars Prevent killing of pods.

```
- name: php-fpm-exporter
  image: hipages/php-fpm_exporter
  ports:
  - containerPort: 
    protocol: TCP
    name: php-fpm-metrics
  resources:
    requests:
      cpu: 100m
      memory: "128Mi"
    limits:
      cpu: 200m
      memory: "512Mi"
```

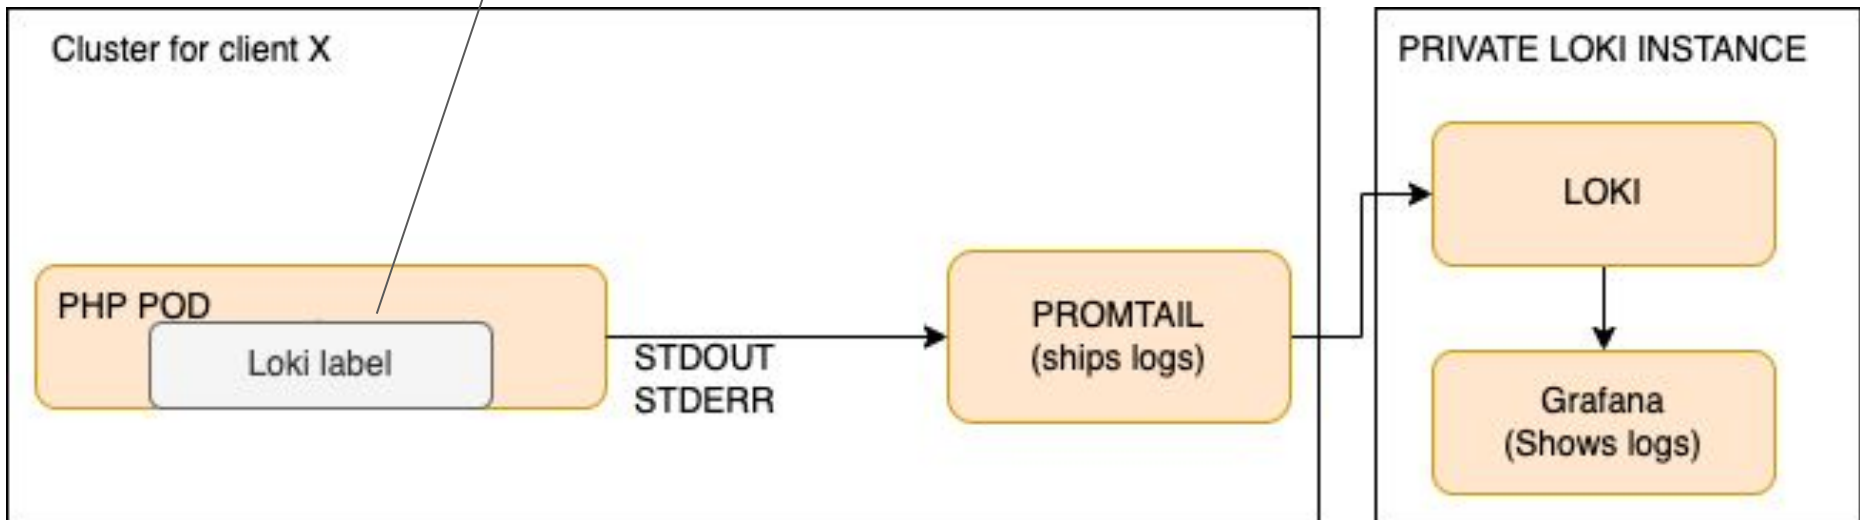
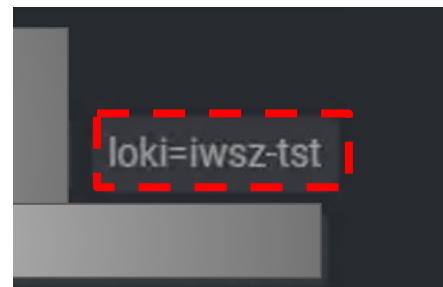
<https://medium.com/finnovate-io/how-to-prevent-kubernetes-cron-jobs-with-sidecar-containers-from-getting-stuck-912c0f1497a3>



# Logs


Labels

```
apiVersion: batch/v1
kind: CronJob
metadata:
  name: "{{.Release.Name}}-cron"
  namespace: {{.Release.Namespace}}
  labels:
    type: "allcronjobs"
    loki: "{{.Values.division }}-{{.Values.environment }}"
```





Infra as code





# Infra as code



## drupal-helm-charts-k8s

This is the representation of the above kubernetes concept specifically for (normal) drupal

```
— CODEOWNERS
— Chart.yaml
— disabled
— readme.md
— templates
  — _helpers.tpl
  — configmap_cypress.yaml
  — configmap_k6.yaml
  — configmap_php_ini.yaml
  — configmap_varnish.yaml
  — cronjob.yaml
  — cronjob_cypress.yaml
  — cronjob_db_dump.yaml
  — cronjob_k6.yaml
  — deployment_maintenance.yaml
  — deployment_nginx.yaml
  — deployment_php.yaml
  — deployment_varnish.yaml
  — hpa_php.yaml
  — deployment_varnish.yaml
  — hpa_php.yaml
  — imagePullSecret.yaml
  — ingress.yaml
  — pod-drush.yaml
  — pvc-dumps.yaml
  — pvc-private.yaml
  — pvc-tmp.yaml
  — pvc.yaml
  — service_maintenance.yaml
  — service_memcached.yaml
  — service_nginx.yaml
  — service_php_fpm.yaml
  — service_varnish.yaml
  — serviceaccount.yaml
  — statefulset_memcached.yaml
— values-env.yaml
```



# Pitfalls




Warning

```
containers:  
- name: {{.Release.Name}}-varnish  
  image: varnish:6.4  
imagePullPolicy: Always
```

## Fix your version

All of a sudden you're running a newer version of varnish in your cluster, this is something you want to test!





# Pitfalls



Warning


```
# Validate helm to not leave a broken state.
```

```
- name: DRY RUN
  run: |
    cd helm
```

```
    helm upgrade --dry-run --timeout 15m --install -n ${{
inputs.app_key }}-${{ inputs.environment }} -f values-${{
inputs.environment }}.yaml --set deploy.runid=${{ github.run_id }}
--set image.tag=${{ inputs.version }} --set image.version=${{
inputs.version }} ${{ inputs.app_key }} .
```

**It's not always correct (sometimes deploys still fail) but will fail if your syntax is wrong in your overrides and will prevent you bringing everything down.**

**Do this before the actual deploy!**





# Pitfalls



Warning

```
- name: DRY RUN
  run: |
    cd helm
    helm upgrade --dry-run --timeout 15m --install -n ${{
inputs.app_key }}-${{{ inputs.environment }} -f values-${{{
inputs.environment }}}.yaml --set deploy.runid=${{{ github.run_id }}}
--set image.tag=${{{ inputs.version }} --set image.version=${{{
inputs.version }} ${{ inputs.app_key }} .
```

```
/**
 * Deployment identifier.
 *
 * Drupal's dependency injection container will be automatically invalidated and
 * rebuilt when the Drupal core version changes. When updating contributed or
 * custom code that changes the container, changing this identifier will also
 * allow the container to be invalidated as soon as code is deployed.
 */
define('DEPLOY_UNIQUE_RUN_ID', $_SERVER['DEPLOY_UNIQUE_RUN_ID'] ?? date('Y-m-d-h'));
$settings['deployment_identifier'] = \Drupal::VERSION . '-' . DEPLOY_UNIQUE_RUN_ID;
```


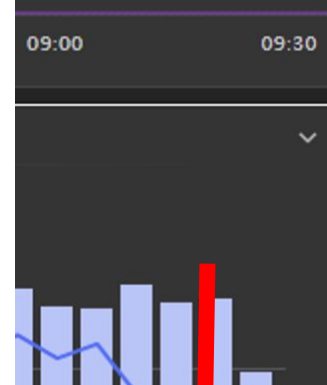
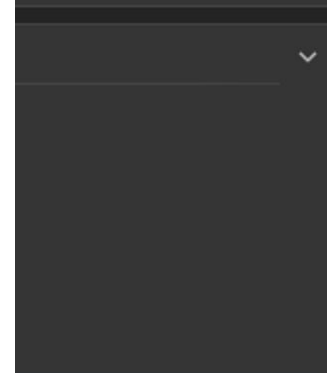



# Pitfalls

The first month in PRD  
The CPU's were throttling, even with not much load on the systems.

Only after checking Dynatrace and tweaking the CPU from 500m to 1 the throttling disappeared (response time decrease & performance increase).

```
resources:  
  requests:  
    cpu: 400m  
    memory: "256Mi"  
  limits:  
    cpu: 1  
    memory: "4112Mi"
```





# Pitfalls



regular POD <VS> Deployment

Use a **Deployment** if you want the pods to be there after a cluster rebuild (eg. global outage).

Obviously we learned this the hard way.

